

A Generic Implementation of the Pruned Dynamic Programming Algorithm

Alice Cleynen ♠

Michel Koskas ♠

Guillem Rigail ♣

alice.cleynen@agroparistech.fr michel.koskas@agroparistech.fr guillem.rigail@genopole.cnrs.fr

♠ UMR 518 AgroParisTech / INRA, 16 rue C. Bernard, 75005, Paris, France

♣ Laboratoire Statistique & Génome, UMR CNRS 8071, INRA,
1152, 523 place des Terrasses, 91000 Evry, France

March 3, 2013

Abstract

We present in this paper a generic implementation of the Pruned Dynamic Programming Algorithm. We discuss the performance of this algorithm compared to that of several algorithms (PELT, CART) - also programed in C++ to allow a fair comparison. The program was written in a full template way, thus allowing a large range of applications and a convenient way of adding extensions.

1 Introduction

We consider a change-point detection framework where we want to obtain the K_{max} best partitions in 1 to K_{max} regions of a signal of size n with respect to some loss or cost function. This problem arises when studying series of data, such as climate data (when did the Maunder minimum, the little ice age or the Dalton minimum begin?) or financial data series. This problem is also frequently encountered in CGH analysis, audio processing and Seq-data analysis. For each $1 \leq k \leq K_{max}$ the signal is divided into k contiguous and homogeneous segments delimited by $k - 1$ splitpoints denoted as $\tau_1, \dots, \tau_{k-1}$. For convenience, we note $\tau_0 = 0$ and $\tau_k = n$. For $1 \leq l \leq k$, a segment r_l is the set of points

$$r_l = \{t \in \mathbb{N} | \tau_{l-1} \leq t < \tau_l\} = \llbracket \tau_{l-1}; \tau_l \rrbracket \quad (1)$$

Difficulties arise both because of the large number of possible partitions ($\binom{n-1}{k-1}$ for a signal of length n divided into k segments) and because many possible optimal solutions may exist.

In the following section we briefly present some algorithms related to our problem, their issues, and why we felt the need for our implementation of the Pruned Dynamic Programming Algorithm. In Section 3 we briefly present the PDPA algorithm, its complexity and the basic operations it relies on as well as its implementation in C++ and a

large range of applications and possible extensions. Finally, In Section 4 we present comparisons of splitpoint positions and runtime between our algorithm and both the DPA and the CART heuristic.

2 Related Work

Several algorithms exist to address segmentation problems. They can be divided into either exact methods, with drawbacks on complexity and capacity, efficient heuristic and constrained-segmentation methods.

The first class essentially consists in improvements of the Dynamic Programing Algorithm (DPA) introduced by ([1]). It finds an optimal solution with respect to the quadratic loss in $\Theta(K_{max}n^2)$ time and $\Theta(n^2)$ space. Improvements include a version that is linear in space ([4]), and the Optimal Partitioning method that directly retrieves an optimal segmentation in an unknown K segments and whose complexity is $\mathcal{O}(n^2)$. Nevertheless two main issues remain: the quadratic complexity in n restricts the use of this algorithm to signals of small length n (reasonable up to $n = 10^4$, too long as soon as $n \geq 10^5$) and the lack of versatility (only the Gaussian loss is implemented) makes it an obstacle to its use for most data-sets (for instance any count data). Recently, the PELT algorithm (Pruned Exact Linear Time [7]) was introduced, that finds the exact optimal segmentation in an unknown number k of segments which is estimated within the algorithm. Even though under certain conditions this algorithm is linear in time ($\mathcal{O}(n)$), PELT only returns the best solution for an unchoosable k .

The second class of algorithms includes a wide range of methods. One of them is the CART algorithm (or Binary Segmentation, [2], [3]), a heuristic procedure to approach the best segmentations in 1 to K_{max} segments. The idea is to split a segment into two segments at each step by minimizing a criterion, and to keep the best partition for the next step. The algorithm can either be stopped when a fixed number of segments is reached, or when the cost gets lower than a fixed threshold. Its complexity is bounded by $\omega(n) \leq C \leq \mathcal{O}(K_{max}n)$ in time and by $\mathcal{O}(k)$ in space (signal data excluded). This method, though faster than the DPA, is not guaranteed to retrieve an optimal solution. Other non-optimal methods include optimization of a constrained criterion. One example is Harchaoui's procedure ([5]) that consists in optimizing the least-square criterion subject to the Lasso constraint $\sum |\mu_t - \mu_{t-1}| \leq s$, where s is a threshold to set.

In the next section we describe the Pruned Dynamic Programing Algorithm (PDPA) proposed in [9] and it's C++ implementation. It represents a considerable improvement to the existing methods as it is exact, faster than the DPA (its time complexity is bounded by $\omega(K_{max}n) \leq C \leq \mathcal{O}(K_{max}n^2)$ and is empirically in $\mathcal{O}(Kn \log(n))$), thrifty (its space complexity is bounded by $\Theta(Kn)$), and highly versatile.

3 The algorithm

3.1 Framework:

Let us assume that we have a sequence of length n : $\{Y_t\}_{t \in \llbracket 1, n \rrbracket}$. We define $\mathcal{M}_{k,t}$ the set of all possible partitions in $k > 0$ regions of the sequence up to point t . The number of possible partitions, $\text{card}(\mathcal{M}_{k,t})$, is $\binom{t-1}{k-1}$. The goal of the pruned DPA is to obtain for all $k \leq K_{\max}$ the contiguous partition m in $\mathcal{M}_{k,t}$ of minimal loss (Lavielle 2005): $\sum_{r \in m} c(r)$, with $c(r)$ the cost of region r of m defined as:

$$\begin{aligned} c(r, \mu) &= \sum_{i \in r} \gamma(Y_i, \mu) + g(\mu) \\ c(r) &= \min_{\mu \in \mathbb{R}} \{c(r, \mu)\}, \end{aligned}$$

where γ is a loss function and g is a regularization penalty.

We define $M_{k,t}$ as the optimal contiguous partition in k regions up to t and $C_{k,t}$ as its optimal cost:

$$\begin{aligned} M_{k,t} &= \arg \min_{\{m \in \mathcal{M}_{k,t}\}} \left\{ \sum_{r \in m} c(r) \right\} \text{ and} \\ C_{k,t} &= \min_{\{m \in \mathcal{M}_{k,t}\}} \left\{ \sum_{r \in m} c(r) \right\}. \end{aligned}$$

As $C_{k,t}$ is region-additive, the following update equation holds:

$$\forall t \geq k \geq 2 \quad C_{k,t} = \min_{k-1 \leq \tau \leq t} \{ C_{k-1,\tau} + c(\llbracket \tau + 1, t \rrbracket) \} \quad (2)$$

Using update equation 2, the original DPA performs t comparisons at each step and retrieves all $\{C_{k,n}\}_{k \in [1, K_{\max}]}$ in exactly $\Theta(K_{\max} n^2)$ runtime.

3.2 Overview of the PDPA

The PDPA studies the function $H_{k,t}(\mu)$ which is the cost of the best contiguous partition in k regions with a last region parameter μ :

$$H_{k,t}(\mu) = \min_{k-1 \leq \tau \leq t} \{ C_{k-1,\tau} + c(\llbracket \tau + 1, t \rrbracket, \mu) \},$$

and from there gets $C_{k,t}$ as $\min_{\mu} \{H_{k,t}(\mu)\}$. More precisely, for each total number of regions, k , from 2 to K , the pruned DPA works on a list of last splitpoint candidates: ListCandidate_k . For each of these candidate splitpoint, τ , the algorithm stores a cost function and a set of optimal-cost intervals. To be more specific, we define:

- $H_{k,t}^{\tau}(\mu) = C_{k,\tau} + \sum_{j=\tau+1}^t \gamma(Y_j, \mu) + g(\mu)$: the optimal cost if the last splitpoint is τ ;
- $S_{k,t}^{\tau} = \{\mu \mid H_{k,t}^{\tau}(\mu) \leq H_{k,t}(\mu)\}$: the set of μ such that τ is optimal;

- $I_{k,t}^\tau = \{\mu \mid H_{k,n}^\tau(\mu) \leq H_{k,n}^t(\mu)\}$: the set of μ such that τ is better than t , with $\tau < t$.

We obviously have $H_{k,t}(\mu) = \min_{\tau \leq t} \{H_{k,t}^\tau(\mu)\}$.

The PDPA rely on four basic properties of these quantities:

1. if all $\sum_{j=\tau+1}^{t+1} \gamma(Y_j, \mu)$ are unimodal in μ then $I_{k,t}^\tau$ are intervals;
2. it is straightforward to get $H_{k,t+1}^\tau(\mu)$ from $H_{k,t}^\tau(\mu)$ using:

$$H_{k,t+1}^\tau(\mu) = H_{k,t}^\tau(\mu) + \gamma(Y_{t+1}, \mu);$$

3. it is easy to update $S_{k,t+1}^\tau$ using:

$$\begin{aligned} S_{k,t+1}^\tau &= S_{k,t}^\tau \cap I_{k,t+1}^\tau \\ S_{k,t}^\tau &= \mathbb{R} \setminus (\cup_{\tau \in \llbracket k-1, t-1 \rrbracket} I_{k,t}^\tau); \end{aligned}$$

4. once it has been determined that $S_{k,t}^\tau$ is empty, τ can be discarded from the list of last splitpoint candidates:

$$S_{k,t}^\tau = \emptyset \Rightarrow \forall t' \geq t \quad S_{k,t'}^\tau = \emptyset.$$

3.3 Complexity of the algorithm

If all $\sum_{j=\tau+1}^{t+1} \gamma(Y_j, \mu)$ are unimodal in μ then the worst case complexity of the algorithm can be bounded in $\Theta(Kn^2)$ time and $\Theta(Kn)$ space (see [9]). In other words the algorithm is at worst equivalent to the original DPA and its improvement by Guédon. This worst case scenario occurs when one has to segment linear data (namely $\forall i, Y_i = \alpha i$ for some $\alpha \in \mathbb{R}$). In that case all positions are kept as possible splitpoints at each step of the algorithm and $O(l)$ comparisons need to be performed at step l , which gives a quadratic time complexity.

However, for many profiles the algorithm is empirically faster. For a constant profile the time complexity is in $\Theta(Kn)$. For such a signal the number of splitpoint candidates at each step is only 1. Empirically, at least for CGH profiles, the number of candidates stored at each step of the algorithm seems to be in $\log(l)$ rather than l and the time complexity is in $\mathcal{O}(Kn \log(n))$.

On a regular computer, the runtime to process a million-point profile and retrieve the best segmentations in 1 to 50 regions is a few minutes. The runtime of the original algorithm is several days.

3.4 General architecture of the package

The main work presented in this paper is the engineering of third author's algorithm. The difficulties we had to come through were the versatility of the program to be written and the design of the objects it could work on. Indeed, the use of full templates implied that we used stable sets of objects for the operations that were to be performed on.

Namely:

- The sets were to be chosen in a *tribe*. This means that they all belong to a set \mathcal{S} of sets such that any set $s \in \mathcal{S}$ can be conveniently handled and stored into the computer and such that

1. if s belongs to s , $\mathbb{R} \setminus s \in \mathcal{S}$
2. if $s_1, s_2 \in \mathcal{S}$, $s_1 \cap s_2 \in \mathcal{S}$
3. if $s_1, s_2 \in \mathcal{S}$, $s_1 \cup s_2 \in \mathcal{S}$

(We call such a set of sets \mathcal{S} an *acceptable* set of sets.)

- The cost functions were chosen in a set \mathcal{F} such that
1. each function may be conveniently handled and stored by the software
 2. for any $f \in \mathcal{F}$, the software can easily solve $f(x) = 0$ and the set of solutions belongs to an acceptable set of sets
 3. for any $f \in \mathcal{F}$ and any constant c , the software can easily solve $f(x) \leq c$ and the set of solutions belongs to an acceptable set of sets
 4. for any $f, g \in \mathcal{F}$, $f + g \in \mathcal{F}$.

This is why we chose the sets of sets to be chosen in collections of intervals or parallelepipeds and implemented the loss functions corresponding to negative binomials, Poisson or normal laws. One can also incorporate a penalty function $g(\mu)$. This function should not depend on the segment characteristic. Typically one can choose $g(\mu) = \lambda|\mu|$ with $\lambda > 0$.

Of course the program is designed in a way that any user could add his own cost function or acceptable set and use it without rewriting a line in the code.

3.5 Link with the algorithm

For each possible number k of regions, from 2 to K_{max} , the pruned DPA proceeds schematically as follows. First, the list of candidates is initialized as $\{k - 1\}$ because the first possible last splitpoint is $k - 1$. Then for every new data point $t \geq k$ the pruned DPA proceeds in five steps:

1. all candidate cost functions ($H_{k,t}^\tau(\mu)$) are compared to the new candidate t to retrieve the interval $I_{k,t}^\tau = \{\mu \mid C_{k-1,\tau} + \sum_{j=\tau+1}^t \gamma(Y_j, \mu) \leq C_{k-1,t}\}$ (property 1);
2. all candidate cost functions ($H_{k,t}^\tau(\mu)$) are updated using property 2;
3. all candidate sets of intervals ($S_{k,t}^\tau$) are updated using the computed $I_{k,t}^\tau$ and property 3;
4. $C_{k,t}$ is computed as $\min_{\text{ListCandidate}_k} \{\min_\mu \{H_{k,t}^\tau(\mu)\}\}$
5. all candidates with an empty $S_{k,t}^\tau$ are discarded using property 4.

To implement these five steps in a versatile fashion we define two template classes, one for the cost function and one for the set of intervals. The template class for a set of intervals has a method to get the complement of a union of intervals and a method to compute its self-intersection with an interval. The template class for the cost function has a $+$ operation (to update the cost function), a method to find its minimum and a method to compute the set of values for which it is smaller than a given constant ($I_{k,t}^T$).

The core of the algorithm is written using these template classes. In this way only a cost function class needs to be specified. As an example and for testing purposes we specified the quadratic, Poisson and negative binomial cost functions. In the quadratic function class the obtention of the interval for which it is minimal is internally computed using analytical formulae whereas in the case of the Poisson and negative binomial this is done using Newton–Raphson’s method.

4 Results and comparisons

4.1 Comparisons of performance

In order to illustrate the validity of our code we have performed a number of simulations for which we have compared the segmentations and costs given by the Original DPA (ODPA) and the Pruned DPA (PDPA). We implemented Guédon’s version of the algorithm in C++ to allow for fair comparisons with many losses (available versions are usually in R or matlab and only for the quadratic loss). A number of possibilities regarding the choice of the size of the signals (varying between 100 and 10 000 points) and number of segments (range of 10 to 150) have been tested for the quadratic, Poisson and negative binomial losses. For each combination, 1000 simulations have been tested, each time assuming the number of segments is given. (For more details regarding the simulation procedure see the Appendix).

In all simulations, our algorithm retrieved the same cost as that given by the ODPA with a tolerated error of up to 10^{-10} to allow for differences of computer approximations in the computations. The percentage of identical segmentations (*i. e.* the number of times both algorithms chose the same splitpoints) for the different losses varies between 99 and 100%. This is due to the possible existence of different equivalent segmentations that are optimal with respect to the log-likelihood criteria. In the particular case of discrete data, this configuration is frequent. We show an example where there are three different and non trivial optimal segmentations in two segments for the quadratic loss (see Figure 1).

We then compared our result to the CART algorithm. Because no C++ version was available, we implemented it for the quadratic loss and only performed the comparisons for the latter.

CART runs the following way: one has an interval, say I , to be split into k pieces but the splits are computed and decided one by one. At the first step one choses the best way to split I in two parts, I_1^1 and I_2^2 such that $I_1^1 \sqcup I_2^1 = I$.

At the second step one has to compute the best splitting of I_1^1 and I_2^1 and keep the best of them. Thus one splits one of the two intervals I_1^1 and I_2^1 and obtains I_1^2 , I_2^2 and I_3^2 and so on...




Figure 1: Three equivalent optimal segmentations with the quadratic loss

The black triangles represent the signal, the colored lines are equivalent optimal segmentations.

Done naively, this leads to a quadratic (in k) number of comparisons. Instead we used a heap that allows for the complexity of CART to be in $O(k \log k)$ instead of the $O(k^2)$ usually observed.

Once again simulations were run for a large combination of signal sizes and segment numbers. Figure 4.1 shows the percentage of correct splitpoints found by CART (black circles) and the relative error between the best cost (obtained by PDPA) and the cost of the segmentation chosen by CART (red triangles) when the size of the signal increases (a.) and when the number of segments increases (b.) We can see that CART performs well when the signal is large and the number of segments is small, but fails to retrieve optimal segmentations when the ratio of the size of the signal over the number of segments gets small. In a total of 2200 simulations, CART retrieved the correct segmentation 43 times.

Let us finally emphasize again the fact that our software works with many possible loss functions or cost functions and sets (see subsection 3.4.)

4.2 Comparison of runtimes

In this section we compared the speed performance of the three algorithms. Because the original dynamic programming algorithm is known to be quadratic in time, the comparison with it was done only for signals up to 100,000 points. For larger signals experiments were done individually. The reader may refer to Figure 3 for a summary of results.

Then we compared the performance of PDPA and CART regarding speed. While the latter is faster, it is important to note that the shape of the curves is identical, as shown in Figure 4.

While implementing CART, one is confronted with the following problem. At step l one has l intervals partitioning the initial interval and has the score of each of them. The interval obtaining the best score needs to be split in two, its two parts need to have their score computed and to be added to the pool of intervals.

The best way to handle this problem is to use a heap and this is what the authors did for the comparisons below.

Let us now compute an approximation of the number of operations needed for both




Figure 2: Comparison of the performance of CART and PDPA

On the left figure the X axis represents the size of the signal on the log-scale, on the right figure the number of segments. For both figures, the left Y axis represents the percentage of right splitpoints retrieved by CART, the right Y axis the relative error of CART when compared to the cost of the optimal segmentation (given by PDPA).




Figure 3: Comparison of the runtimes of ODPA and PDPA

Both figures compare the speed of OPDA and PDPA for signals with varying sizes (given on the X axis) but a fixed number of segments (15). The results are given in seconds (Y axis).




Figure 4: Comparison of the runtimes of CART and PDPA

On the left figure the X axis represents the size of the signal on the log-scale, on the right figure the number of segments. For both figures, the left Y axis represents the speed of the PDPA while the right Y axis represents the speed of CART, both in seconds. Even though CART is a lot faster, the shape of the curves are similar.

algorithms. At each new step, in order to find the best division of a segment of length l into 2 pieces, CART requires $7l - 3$ elementary operations. Thus the worst number of elementary operations is bounded by $K_{max}(7n - 3) + \sum_1^{K_{max}} \log_2(k)$. (The term $\log_2(k)$ comes from the use of the heap.)

If the splits are quite regularly distributed, one obtains $(7n - 3) \log_2(K_{max}) + \sum_1^{K_{max}} \log_2(k)$. Therefore one may expect a $n \log_2(K_{max})$ time complexity.

Now PDPA needs at each step to find the roots and minimum of cost functions for each candidate, for which it requires 13 elementary operations. Intersecting two intervals requires 2 elementary operations. If the pruning is perfect (one candidate left at each step), then PDPA will require $15K_{max}n$ elementary operations. Therefore the best configuration for PDPA is still worse than the worst configuration for CART.

Of course one should keep in mind that CART does not find the best segmentations but in most cases quickly finds quite good segmentations.

In our simulation study we chose the number of segments and the size of the signal in order to analyse the general trend of the algorithms and to mimic signals expected in many applications: for instance, a segmentation in 15 segments is common in CGH-array, a length of 10^6 or 10^7 is common with Seq-dat, etc.

Recently, ([6],2012) compared the ability of several segmentation approaches to segment real Comparative Genomic Hybridization data. They show in table 3 (see line cghseg.k) that the PDPA (we implemented a preliminary version of it in the cghseg package) coupled with an appropriate model selection scheme outperforms existing method in terms of error. This demonstrate the advantage of the PDPA on real CGH data.

5 Discussion and Conclusion

We described a generic C++ implementation of the PDPA described in [9] and that gets the best partitions in 1 to K_{max} regions for a given loss function. We made an extensive comparison of the result of our code with the one obtained with the ODPa to illustrate its validity, and we implemented several loss functions (including the frequently used quadratic one) in our code to demonstrate its versatility. A user who might be interested in a particular loss function adapted to his data would only need to give its expression, and formulae to compute its minimum and roots (and is welcome to write to one of the authors for help). So far the implementation is limited to a set of one-parameter loss functions, but future work should include an extension to multi-parameter losses.

The comparison of the result of the PDPA to the CART heuristic proves the advantage of retrieving the best partitions and a preliminary implementation of our code was recently shown to provide state of the art results on real CGH data (Hocking et al 2012). Moreover, this implementation can be used as a base for further analysis. For example, [8] use it to initialize their Hidden Markov Model to compute change-point location probabilities.

Our code and an R package including our C++ code are available on the CRAN (or send an e-mail to one of the authors).

APPENDIX

The signals used for the comparisons between CART, ODPa and PDPA were generated as described in this paragraph. Once the size of the signal, the number of segments and the generating function (normal, Poisson or negative binomial laws) were chosen, the splitpoints were equally spaced along the signal and the parameters chosen as follows:

- for the normal law, the variance parameter σ^2 was set to 1 along the sequence. Then for each segment k , the mean parameter μ was chosen equal to $\mu_k = 2r$ with $r = k - 6q$ and $1 \leq r \leq 6$.
- Similarly, the parameter of the Poisson was subject to the same pattern ($\lambda_k = 2r$ with $r = k - 6q$ and $1 \leq r \leq 6$).
- For the negative binomial law we used the parametrization $\mathcal{NB}(\mu, \phi)$ where $\phi = \mathbb{E}(X)^2 / (\text{Var}(X) - \mathbb{E}(X))$ and $\mu = (\text{Var}(X) - \mathbb{E}(X)) / \text{Var}(X)$. For the simulations, ϕ was fixed equal to 0.6 for the whole sequence and we selected $\mu_k = 1/(r + 1)$ with $r = k - 6q$ and $0 \leq r < 6$.

References

- [1] R. Bellman. On the approximation of curves by line segments using dynamic programming. *Commun. ACM*, 4(6):284, 1961.
- [2] L. Breiman, J. H. Friedman, R. Olshen, and C. J. Stone. Classification and regression trees. *Wadsworth and Brooks*, 1984.

- [3] S. Gey and E. Lebarbier. Using CART to detect multiple change points in the mean for large sample. http://hal.archives-ouvertes.fr/hal-00327146_v1/, February 2008.
- [4] Y. Guédon. Exploring the segmentation space for the assessment of multiple change-point models. Technical Report 6619, INRIA, 2008.
- [5] Z. Harchaoui and C. Levy-Leduc. Catching change-points with lasso. In J.C. Platt, D. Koller, Y. Singer, and S. Roweis, editors, *Advances in Neural Information Processing Systems 20*, pages 617–624. MIT Press, Cambridge, MA, 2008.
- [6] T. D. Hocking, G. Schleiermacher, I. Janoueix-Lerosey, O. Delattre, F. Bach, and J.-P. Vert. HAL - INRIA :: [hal-00663790, version 1] learning smoothing models of copy number profiles using breakpoint annotations. <http://hal.inria.fr/hal-00663790>.
- [7] R. Killick, P. Fearnhead, and I. Eckley. Optimal detection of changepoints with a linear computational cost. *Arxiv:1101.1438v2*, April 2011.
- [8] T.M. Luong, Y. Rozenholc, and G. Nuel. Fast estimation of posterior probabilities in change-point models through a constrained hidden markov model. *Arxiv preprint arXiv:1203.4394*, 2012.
- [9] G. Rigai. Pruned dynamic programming for optimal multiple change-point detection. *Arxiv:1004.0887*, April 2010.